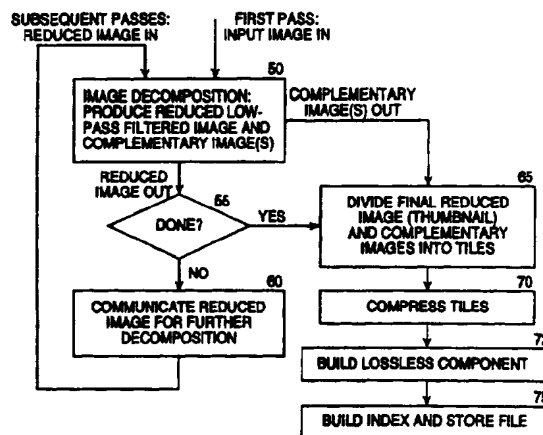




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06K 9/36, 9/42, 9/54, G06F 15/00, H04N 7/12, 11/04	A1	(11) International Publication Number: WO 97/17669 (43) International Publication Date: 15 May 1997 (15.05.97)
(21) International Application Number: PCT/US96/18017 (22) International Filing Date: 7 November 1996 (07.11.96) (30) Priority Data: 08/554,384 8 November 1995 (08.11.95) US (60) Parent Application or Grant (63) Related by Continuation US 08/554,384 (CIP) Filed on 8 November 1995 (08.11.95) (71) Applicant (for all designated States except US): STORM TECHNOLOGY, INC. [US/US]; 521 Almanor Avenue, Sunnyvale, CA 94086 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): STARREVELD, Adolf, G. [NL/US]; 1555 West Middlefield Road #111, Mountain View, CA 94043 (US). LIGTENBERG, Adrianus [NL/US]; 735 Holly Oak Road, Palo Alto, CA 94303 (US). GUL- LAND, William, J. [GB/US]; Apartment 51, 1375 Mon- tecity Avenue, Mountain View, CA 94043 (US).		(74) Agents: SLONE, David, N. et al.; Townsend and Townsend and Crew L.L.P., 8th floor, Two Embarcadero Center, San Francisco, CA 94111-3834 (US). (81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, HU, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, TJ, TM, TR, TT, UA, UG, US, UZ, VN, ARIPO patent (KE, LS, MW, SD, SZ, UG), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, ML, MR, NE, SN, TD, TG). Published <i>With international search report.</i>

(54) Title: METHOD AND FORMAT FOR STORING AND SELECTIVELY RETRIEVING IMAGE DATA

**(57) Abstract**

A method of processing an input image for storage includes decomposing (50) the input image into a number of images at various resolutions, subdividing (65) at least some of these images into tiles (rectangular arrays) and storing (75) a block (referred to as the "tile block") representing each of the tiles, along with an index that specifies the respective locations of the tile blocks. In specific embodiments, the tiles are 64 x 64 pixels or 128 x 128 pixels. The representations of the tiles are typically compressed versions (78) of the tiles. In a specific embodiment, JPEG compression is used. In a specific embodiment, an operand image is recursively decomposed to produce a reduced image and a set of additional (or complementary) pixel data. At the first stage, the operand image is normally the input image, and, for each subsequent stage, the operand image is the reduced image from the previous stage. At each stage, the reduced image is at a characteristic resolution that is lower than the resolution of the operand image. The processing is typically carried out until the resulting reduced image is of a desired small size (55).

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AM	Armenia	GB	United Kingdom	MW	Malawi
AT	Austria	GE	Georgia	MX	Mexico
AU	Australia	GN	Guinea	NE	Niger
BB	Barbados	GR	Greece	NL	Netherlands
BE	Belgium	HU	Hungary	NO	Norway
BF	Burkina Faso	IE	Ireland	NZ	New Zealand
BG	Bulgaria	IT	Italy	PL	Poland
BJ	Benin	JP	Japan	PT	Portugal
BR	Brazil	KE	Kenya	RO	Romania
BY	Belarus	KG	Kyrgyzstan	RU	Russian Federation
CA	Canada	KP	Democratic People's Republic of Korea	SD	Sudan
CF	Central African Republic	KR	Republic of Korea	SE	Sweden
CG	Congo	KZ	Kazakhstan	SG	Singapore
CH	Switzerland	LI	Liechtenstein	SI	Slovenia
CI	Côte d'Ivoire	LK	Sri Lanka	SK	Slovakia
CM	Cameroon	LR	Liberia	SN	Senegal
CN	China	LT	Lithuania	SZ	Swaziland
CS	Czechoslovakia	LU	Luxembourg	TD	Chad
CZ	Czech Republic	LV	Latvia	TG	Togo
DE	Germany	MC	Monaco	TJ	Tajikistan
DK	Denmark	MD	Republic of Moldova	TT	Trinidad and Tobago
EE	Estonia	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	UG	Uganda
FI	Finland	MN	Mongolia	US	United States of America
FR	France	MR	Mauritania	UZ	Uzbekistan
GA	Gabon			VN	Viet Nam

**METHOD AND FORMAT FOR STORING AND SELECTIVELY
RETRIEVING IMAGE DATA**

5

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner
10 has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

15

BACKGROUND OF THE INVENTION

The present invention relates generally to storing and retrieving image data, and more specifically to a file format that provides selective retrieval for display, printing, or communication purposes.

20

It is said that a picture is worth a thousand words, and history bears this out. From cave men drawing pictures on the walls of their caves to modern-day students accessing multimedia encyclopedias, people have always considered pictorial information an essential communication tool. Recent advances in computer technology have made it
25 possible to create and exchange elaborate documents in electronic form. Text and graphics (line art) are easily formatted and manipulated, allowing an amateur computer user to create highly professional documents. The logical next step is the incorporation of photographs and other bitmap images into documents.

25

30

While the technology exists to digitize and edit photographs, a persistent problem is that bitmap images are very large. For example, an 8x10 inch color photograph scanned at 300 dpi (dots/inch) at 24 bits/pixel represents over 20 megabytes of data, and this is hardly an extreme example. Thus, in the context of a local computer, acquiring,
35 viewing, manipulating, and printing such images are painfully slow processes, even with state-of-the-art computer equipment. The problem is aggravated when it is desired to transfer these images over a network.

35

Fortunately, the bitmap representation of most photographs includes a large amount of repetitive information making the images
40 amenable to image compression schemes. An international compression standard called JPEG, which allows pictures to be stored in about 1/10 the space with high quality, has been developed and adopted. Greater levels of compression can be achieved, but the quality of the reconstructed image is lower. While it is possible to achieve lossless compression, the
45 compression factor would only be on the order of 4x. A discussion of the JPEG compression, as well as a number of enhancements thereto, can be

40

45

found in U.S. Patent No. 5,333,212 to Ligtenberg, the entire disclosure (including appendices) of which is incorporated by reference for all purposes.

There have been developed image editing programs that incorporate virtual memory schemes specially tailored to images. In these schemes, the images are divided into tiles (rectangular image regions). The advantage of tiles is that image data can be retrieved for only those tiles that will be displayed, and operations can be performed locally.

Another problem is where an image needs to be accessed at different resolutions. For example, an image might be displayed actual size at 72 dpi but printed at 300 or 600 dpi. A solution to this problem is known as pyramid coding, such as in Kodak's PhotoCD format, where the image is stored at different resolutions. The pyramid allows the user to select an image resolution that is the most effective for a certain task. For example, to browse a number of images, one views the thumbnail images (highly reduced images that are on the order of an inch or two on a side when displayed). The basic features of the image can be discerned, and a selected image can then be retrieved at a screen resolution for viewing or at a higher print resolution suitable for printing.

Although the above solutions have been successful in addressing some of the obstacles standing in the way of using images on the computer, the solutions are not without their drawbacks. For example, JPEG compression, while it reduces the file size, does not allow selective reconstruction of a desired portion of the image. Rather, it is necessary to reconstruct all portions of the image between the top left of the image and the bottom right of the desired portion. Furthermore, while the tile-based virtual memory schemes alleviate the memory requirements, they do not reduce the file size. Additionally, an image stored in the PhotoCD format is so large that it is only used on PhotoCDs. The format is not practical as a format for on-line storage and retrieval or for storage on a local hard disk.

SUMMARY OF THE INVENTION

The present invention provides an effective and flexible image data format and techniques for selectively storing and retrieving image data. A storage format according to the invention allows desired portions of the image data to be retrieved at desired resolution.

In brief, a method of processing an input image for storage includes decomposing the input image into a number of images at various resolutions, subdividing at least some of these images into tiles (rectangular arrays) and storing a block (referred to as the "tile block") representing each of the tiles, along with an index that specifies the respective locations of the tile blocks. The representations of the tiles are typically compressed versions of the tiles. In a specific embodiment, JPEG compression is used.

In a specific embodiment, an operand image is iteratively decomposed to produce a reduced image at a lower resolution and a set of additional (or complementary) pixel data. At the first stage, the operand image is normally the input image, and, for each subsequent stage, the operand image is the reduced image from the previous stage. The additional pixel data produced at a given stage can be considered to have the same characteristic resolution as the reduced image produced at that stage. The reduced image and the additional pixel data at a given resolution are sometimes referred to as a layer.

The processing is typically carried out until the resulting reduced image is of a desired small size. This is typically a size that represents the smallest image that is useful for display. In a specific embodiment, this is an image between 80 and 160 pixels on a side, referred to as a thumbnail.

The reduction performed at each stage is of a type that retains the basic appearance of the operand image. In a specific embodiment, the reduced image is obtained by subjecting the operand image to horizontal and vertical low-pass filters followed by decimation. In a specific embodiment, the reduction is by a factor of 2 in each dimension.

The additional pixel data contains at least some of the information lost in the reduction of the operand image. This allows the operand image to be reconstructed to a desired degree by suitably combining the reduced image and the additional pixel data. If the additional pixel data allows faithful reconstruction of the operand image, the reduced images (other than the thumbnail) are redundant, and it is not necessary to store them in the file. That is, a given layer may only need to contain the additional pixel data at that layer's resolution. However, in some embodiments, the reduced images are included in the file to allow direct access and to eliminate or reduce the computational resource needed to reconstruct the reduced image.

In a specific embodiment, the additional pixel data comprises a set of complementary images, which are produced by subjecting the operand image to each of three series of filtering operations: a horizontal low-pass filter and a vertical high-pass filter; a horizontal high-pass filter and a vertical low-pass filter; and a horizontal high-pass filter and a vertical high-pass filter. Each of the three resulting filtered images is decimated.

Any desired portion of the image file, can be retrieved and reconstructed at a desired one of the resolutions that characterize the reduced images that were generated during the decomposition process. If a particular reduced image is stored in the file, its relevant tile blocks can be retrieved directly. If the final reduced image is the only reduced image actually stored in the file, the reduced image will typically need to be reconstructed. This is done, in essence, by reversing the decomposition process used to create the file.

In a specific embodiment, the relevant tile blocks of the final reduced image and of the complementary images at the final stage are retrieved, decompressed (if necessary), upsampled, and combined. This produces an expanded image that is an approximation to the reduced
5 (operand) image that was input to the last stage of the decomposition process. The faithfulness of the reconstruction depends on the lossiness of the compression that was done on the tiles, and on the degree to which the complementary image process preserves pixel values.

This process is iterated by combining the reconstructed image
10 with the relevant portions of the complementary images having the same resolution as the reconstructed image. The iterations stop when an image of the desired resolution is achieved. This may be the original input image, or an image at some intermediate resolution. If a lossy compression scheme is used, it is also possible to store in the file
15 additional tiled information that represents the difference between the input image, and the reconstructed image at the input image's resolution. This information could be losslessly compressed.

The present invention also provides efficient techniques for modifying images in memory and updating the file to incorporate the
20 modifications. When a portion of the image is retrieved from the file and brought into memory, a table is created in memory containing index information regarding the tile blocks in the file. The table includes a "valid" bit for each tile block. The valid bit is initially set to a state designating that the tile has not been modified in memory and
25 therefore corresponds to what is currently in the file.

If a portion of the image is modified, the tile blocks that include the modified portion are marked "invalid," signifying that the mapped tiles stored in the disk file no longer correspond to what is in memory. When it is desired to save the modification, the tile blocks for
30 the modified tiles are written to the file, preferably at the end of the file, and the index in the file is modified to reflect the new locations of the modified tile blocks.

A further understanding of the nature and advantages of the present invention may be realized by reference to the remaining portions
35 of the specification and the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a system in which the present invention may be embodied;

40 Fig. 2 is a high-level flow diagram showing the conversion of an input image to a file according to the present invention;

Fig. 3 is an expanded flow diagram illustrating steps iteratively performed during image decomposition;

45 Fig. 4 is a high-level flow diagram showing the retrieval and reconstruction of an image stored in a file according to the present invention; and

Fig. 5 is an expanded flow diagram illustrating steps iteratively performed during image reconstruction.

System Overview

5 Fig. 1 is a simplified block diagram of a computer system 10 in which the present invention may be embodied. The computer system configuration illustrated at this high level is standard, and as such, Fig. 1 is labeled "Prior Art." A computer system such as system 10, suitably programmed to embody the present invention, however, is not prior
10 art. The specific embodiments of the invention are embodied in a general-purpose computer system such as shown in Fig. 1, and the remaining description will generally assume that environment. However, the invention may be embodied in dedicated photo input and output devices such as digital cameras, set-top boxes, and printers.

15 In accordance with known practice, the computer system includes a processor 12 that communicates with a number of peripheral devices via a bus subsystem 15. These peripheral devices typically include a memory subsystem 17, a user input facility 20, a display subsystem 22, output devices such as a printer 23, and a file storage
20 system 25.

In this context, the term "bus subsystem" is used generically so as to include any mechanism for letting the various components of the system communicate with each other as intended. With the exception of the input devices and the display, the other components need not be at the
25 same physical location. Thus, for example, portions of the file storage system could be connected via various local-area or wide-area network media, including telephone lines. Similarly, the input devices and display need not be at the same location as the processor, although it is anticipated that the present invention will most often be implemented in
30 the context of PCs and workstations.

Bus subsystem 15 is shown schematically as a single bus, but a typical system has a number of buses such as a local bus and one or more expansion buses (e.g., ADB, SCSI, ISA, EISA, MCA, NuBus, or PCI), as well as serial and parallel ports. Network connections are usually established
35 through a device such as a network adapter on one of these expansion buses or a modem on a serial port. The computer system may be a desktop system or a portable system.

Memory subsystem 17 includes a number of memories including a main random access memory (RAM) 30 and a read only memory (ROM) 32 in
40 which fixed instructions are stored. In the case of Macintosh-compatible personal computers this would include portions of the operating system; in the case of IBM-compatible personal computers, this would include the BIOS (basic input/output system).

User input facility 20 typically includes a keyboard 40 and
45 may further include a pointing device 42 and a scanner 43. The pointing device may be an indirect pointing device such as a mouse, trackball,

touchpad, or graphics tablet, or a direct pointing device such as a touchscreen incorporated into the display.

Display subsystem 22 typically includes a display controller 44 and a display device 45 coupled to the controller. The display device may be a cathode ray tube (CRT), a flat-panel device such as a liquid crystal display (LCD), or a projection device. Display controller provides control signals to the display device and normally includes a display memory (not shown in the figure) for storing the pixels that appear on the display device.

The file storage system provides persistent (non-volatile) storage for program and data files, and typically includes at least one hard disk drive 46 and at least one floppy disk drive 47. There may also be other devices such as a CD-ROM drive 48 and optical drives. Additionally, the system may include hard drives of the type with removable media cartridges. As noted above, one or more of the drives may be located at a remote location, such as in a server on a local area network or at a site on the Internet's World Wide Web.

File Creation and Format

Fig. 2 is an overall flow diagram showing the conversion of an input image to a file having the file format according to an embodiment of the present invention. The file format stores sufficient information to allow access to desired portions of an original bitmap input image at a variety of resolutions, depending on the task at hand. The file format is referred to as an active file format since the user can interact with the file and the file provides what the user needs.

In the specific embodiment described below, the file is constructed with a view to eliminating redundant information so as to produce a smaller file. The trade-off is that it is necessary to reconstruct the image at a given resolution on the basis of information retrieved from the file.

The process of constructing the file is iterative in the sense that in a first pass the input image is subjected to an image decomposition stage 50, and in subsequent passes, a portion of the data resulting from stage 50 is again input to stage 50. The input image to stage 50 is referred to generically as the operand image. The operation of stage 50 is to produce a reduced image at a lower resolution than the operand image, and additional image components containing a desired level of the information that was discarded in producing the reduced image. These additional image components are referred to as additional pixel data or as the "complementary images." In a specific embodiment, the operations in stage 50 include filtering and decimation by a factor of two.

The reduced image and complementary images produced at each stage have a characteristic resolution and can be thought of as a layer of the image. In the specific embodiment, the reduced image is not stored as

part of the file since it can be reconstructed from information in subsequently generated layers.

When the operand image has been processed, resulting in a reduced image and complementary images, the complementary images are output from stage 50 for further processing and storage as will be described below. The system tests (step 55) whether there is need for another iteration. The method can be iterated any desired number of times, but in a specific embodiment, the test is whether the reduced image has a size below a threshold, which in the specific embodiment is that each dimension of the reduced image be between 80 and 160 pixels. Clearly, the necessary number of iterations can be readily determined ahead of time from the size of the input image.

If the reduced image is not sufficiently small, the reduced image is communicated to stage 50 as the operand image (step 60). If the reduced image is sufficiently small, further processing by stage 50 is not required. At this point, the reduced image, referred to as the "final reduced image" or the "thumbnail" is output for further processing and storage along with the complementary images, which were output at each stage.

The thumbnail and the complementary images (and the other reduced images if they are to be stored as part of the file) are divided into tiles (step 65) to allow subsequent retrieval and reconstruction of desired portions of the image at desired resolutions. The individual tiles are compressed (step 70), although compression is not a necessary part of the invention. Compression, if it is used, is one possible mechanism that can cause loss of information. Although there are lossless compression techniques, the specific embodiment uses JPEG compression, which is a lossy technique.

Another possible source of information loss is the processing within stage 50. The reduced image output from stage 50, by definition, contains only a portion of the operand image that was input to stage 50. The degree to which the operand image can be reconstructed depends on the nature of the complementary image generation. The complementary images may contain all the information to allow such reconstruction if desired, but there may be some applications where there is no need to reconstruct any of the intermediate resolution images losslessly. It may, however, be desired to be able to reconstruct the original input image losslessly.

If it is required to provide lossless reconstruction, the compressed tiles are subjected to processing (step 72) to build a lossless component for storage in the file. This may entail reconstruction of the image at full resolution, subtraction from the original input image, and generation of data representing the difference image. The latter would be tiled and losslessly compressed.

The file comprises individual blocks for each of the tiles produced at tiling step 65, as well as an index to all these blocks. If compression is used, the block would store the compressed data produced at

compression step 70. These blocks are sometimes referred to as "tile blocks." If step 72 is performed, the index would also refer to the lossless information component. The construction of the index and storing of the file is shown as step 75, performed after all the processing described above.

As mentioned above, the file includes an index that contains the locations (offsets in file) of all the tile blocks. In embodiments where each tile is compressed, the tile blocks are of unequal length, and so it is also necessary to store the length of each tile block as well as the tile block's offset in the file. The lengths of the tile blocks can be stored in the index or as part of the tile blocks themselves. Where compression is not used, or where a compression regime that compresses to a fixed size is used (see, for example, the above mentioned U.S. Patent No. 5,333,212), the tile blocks are all of equal length in a given layer, and there is no need to store the individual lengths of the tile blocks. If, as in some alternative embodiments, it is desired to store one or more of the intermediate reduced images themselves, the index would need to store the locations of the tile blocks for these reduced images.

The choice of tile size is subject to a number of design considerations. Since a given image is unlikely to have dimensions that are exact multiples of the tile dimension, smaller tiles will reduce wasted storage of blank space along the bottom and right edges of the image. Further, since users will generally be unaware of tile boundaries when specifying regions for reconstruction, smaller tiles will reduce the amount of unwanted image portions that are reconstructed. On the other hand, larger tiles reduce the number of I/O operations for retrieval and storage. Larger tiles also result in a smaller index, which is not a major issue for storage, but can be significant for memory since the entire index is likely to be in memory.

In general, the number of pixels on a side should be a power of 2 for ease in computing and associating corresponding tiles in different layers (assuming a 2x reduction at each stage), and square tiles are generally preferred. There is no fundamental reason why the tile size should be the same for all layers, and in some embodiments it may be advantageous to have the tiles progressively smaller so that a single tile in one layer will correspond to a single (smaller) tile in the next reduced layer. Furthermore, a particular size may optimize I/O or memory operations for a frequently used layer. However in most of the discussion that follows, a fixed tile size will be assumed. In a current implementation, a tile size of 128x128 pixels is used, but a tile size of 64x64 pixels is another likely choice.

Depending on the tile size and the number of stages of reduction, the reduced image and its complementary images at some stage (for example, the thumbnail stage) may fit within a single tile. In such a case, reference to dividing a particular image into tiles should be interpreted to include the possibility that the image has been

sufficiently reduced so that it fits within a single tile. The benefits of tiling will be realized, however, as long as at least some of the layers are divided into a plurality of tiles.

5 The above description of the method of processing operand images in terms of producing a number of resulting images, tiling and compressing the resulting images, building the index of tile blocks, and storing the file accurately represents the operations that are performed. It should be realized, however, that this is a high-level view of the method. In actual implementations, the steps are likely to be subdivided and interleaved, producing compressed tiles and corresponding index
10 information incrementally as processing proceeds.

One possible approach is to process portions of the input image to completion before processing other portions. For example, if there are to be two levels of reduction, 16 tiles of the input image map
15 to a single tile of the thumbnail and each group of 16 tiles would be processed fully as follows. In the first stage, each of four subgroups of four tiles from the input image produces a single reduced tile and three complementary tiles, the latter of which can be compressed, indexed, and stored immediately. The four tiles of the reduced image produce a single
20 tile of the thumbnail and three complementary tiles, all of which can be compressed, indexed, and stored immediately. At this point, the compressed tiles from the two stages of processing can be used to reconstruct 16 tiles at the resolution of the input image, and the lossless component for the 16 input tiles can be constructed, compressed
25 losslessly, indexed, and stored.

Fig. 3 is an expanded flow diagram illustrating steps performed in a specific embodiment of image decomposition stage 50. In brief, the reduced image and the three complementary images are obtained by subjecting the operand image to four combinations of filtering and
30 decimation operations. The technique, outlined below, is described for use in a different context in U.S. Patent No. 4,943,855 to Bheda et al., the entire disclosure of which is incorporated by reference for all purposes.

In the specific embodiment, the operand image is subjected to
35 a horizontal low-pass filter applied to each row (step 80) followed by a decimation of every other column (step 82), to produce a first intermediate image. The operand image is separately subjected to a horizontal high-pass filter applied to each row (step 85) followed by a decimation of every other column (step 87), to produce a second
40 intermediate image. Each intermediate image is then subjected separately to two additional sets of operations. In the specific embodiment, the filters are finite impulse response (FIR) filters. For example, the low-pass filters may have coefficients (1, 1) or (1, 3, 3, 1) while the high-pass filters may have coefficients (1, -1) or (1, 3, -3, -1).

45 The first intermediate image is subjected to a vertical low-pass filter applied to each column (step 90) followed by a decimation of

every other row (step 92), to produce the reduced image, also referred to as the low-low image. The first intermediate image is separately subjected to a vertical high-pass filter applied to each column (step 95) followed by a decimation of every other row (step 97), to produce a first complementary image, also referred to as the low-high image.

The second intermediate image is subjected to a vertical low-pass filter applied to each column (step 100) followed by a decimation of every other row (step 102), to produce a second complementary image, also referred to as the high-low image. The second intermediate image is separately subjected to a vertical high-pass filter applied to each column (step 105) followed by a decimation of every other row (step 107), to produce a third complementary image, also referred to as the high-high image.

Image Reconstruction

Fig. 4 is an overall flow diagram showing the retrieval and reconstruction of an image from a file having the file format described above. Since the thumbnail and the complementary images at each stage in the image decomposition were tiled, it is possible to selectively reconstruct a desired portion of the image. While the discussion that follows will typically refer to retrieving a particular image, it should be understood that it is only necessary to retrieve the tile blocks and reconstruct the tiles corresponding to the desired portion of the image.

The reconstruction process is, in large part, a reverse of the decomposition process described above. Again, the process is iterative. In a first pass the thumbnail image and its complementary images are retrieved, decompressed (if previously compressed), and subjected to an image expansion and reconstruction stage 120; in subsequent passes, the expanded image resulting from stage 120 is again input to stage 120, along with the retrieved complementary images that had been generated from the image at the next higher level of resolution. The thumbnail image and the expanded image from the previous stage that is again input to stage 120 are referred to generically as the operand image. The operation of stage 120 is to produce an expanded image at a higher resolution than that of the operand image and its complementary images. In a specific embodiment, the operations in stage 120 include interpolation, filtering, and combination of the operand images.

When the operand image and complementary images have been processed, resulting in an expanded image, the system tests (step 125) whether there is need for another iteration. The method can be iterated any desired number of times up to the maximum resolution, namely that of the original input image.

If the expanded image is not yet at the desired resolution, the expanded image is communicated to stage 120 as the operand image (step 130) along with its complementary images. If the expanded image is at the desired resolution, further processing by stage 120 is not required. At

this point, the expanded image, referred to as the final expanded image is output for its intended use.

The above description was in terms of retrieving and reconstructing the image from the file. As mentioned, this entails retrieving the thumbnail and various of the complementary images into memory and, if the tile representations are compressed, decompressing them before operating on them. It should be noted that it is not always necessary to start from the thumbnail. If the desired portion of the image is already in memory at some resolution, either having been reconstructed or directly retrieved, that image can be considered the operand image, which can be expanded and combined with its complementary images. For embodiments where the intermediate reduced images are not stored, the portion that was already in memory is most likely to have been reconstructed starting with the thumbnail and its complementary images.

Fig. 5 is an expanded flow diagram illustrating steps performed in a specific embodiment of image reconstruction stage 120. This embodiment assumes a specific decomposition as shown in Fig. 3. In brief, the expanded image is obtained by subjecting the operand image and the three complementary images to four combinations of upsampling, filtering, and summing operations. The technique, outlined below, is also described in the above-mentioned U.S. Patent No. 4,943,855.

In the specific embodiment, the operand image is vertically upsampled (step 140) and each column is subjected to a vertical low-pass filter (step 142) to produce a first intermediate image. The low-high image is vertically upsampled (step 145) and each column is subjected to a vertical high-pass filter (step 147) to produce a second intermediate image. The first and second intermediate images are then summed (step 150) to produce a third intermediate image. The third intermediate image is horizontally upsampled (step 152) and each row is subjected to a horizontal low-pass filter (step 155) to produce a fourth intermediate image.

The high-low image is vertically upsampled (step 160) and each column is subjected to a vertical low-pass filter (step 162) to produce a fifth intermediate image. The high-high image is vertically upsampled (step 165) and each column is subjected to a vertical high-pass filter (step 167) to produce a sixth intermediate image. The fifth and sixth intermediate images are then summed (step 170) to produce a seventh intermediate image. The seventh intermediate image is horizontally upsampled (step 172) and each row is subjected to a horizontal high-pass filter to produce an eighth intermediate image.

The fourth and eighth intermediate images are summed (step 177) to provide the expanded image that serves as the operand image for the next stage, if further iterations are needed. In the specific embodiment, the filters are finite impulse response (FIR) filters. For example, the low-pass filters may have coefficients (1, 1) or (1, -3, -3,

1) while the high-pass filters may have coefficients $(-1, 1)$ or $(-1, 3, -3, 1)$, assuming decomposition with the specific filters discussed above.

Retrieval at Increased Resolution

5 As described above, with reference to Figs. 4 and 5, retrieving the image at a given resolution entails iteratively combining the operand image with complementary images. While these operations are being carried out in memory, it is not necessary to delay displaying the image. As soon as an image at a given resolution is available, it can be
10 expanded to a size suitable for display using an interpolation such as a cubic spline, and the interpolated values can be sent to the display memory. This provides a display at full size, albeit lower true resolution, without waiting to retrieve the complementary images from disk.

15 Once the complementary images have been retrieved from disk and properly incorporated into the image, the expanded image has the true resolution for its size. If this size is suitable for display, the image can be copied to the display memory. Otherwise, it can be expanded and interpolated as above before overwriting the previous interpolated image
20 in display memory. Since the newly displayed interpolated image includes the complementary images from the last iteration, it represents an improvement over the previously interpolated image. Thus a fuzzy version of the image (the interpolated thumbnail) initially appears on the screen, and sharpens as additional data are retrieved from the disk file and
25 incorporated into the image.

In-Memory Operations and File Saves

 The description of creating and storing a file, as illustrated in Fig. 2, was in terms of starting with an image and iteratively
30 decomposing the image into layers, each including a reduced image and complementary images. While all the complementary images are stored, only the final reduced image is stored. When the image is retrieved and brought into memory at a given resolution, it is only necessary to retrieve the index and the relevant tile blocks to create the appropriate
35 data structures in memory.

 If the image is to be modified, the modifications must be performed on the image at full resolution, even though the image is being displayed at a lower resolution. If it is desired to save the
40 modifications, the tile blocks that include the modified portions must be rewritten to the disk file. In order to make the process more efficient, the data structure corresponding to the index of tile blocks is defined to include a "valid" bit for each tile block. If a portion of the image is modified, the tile blocks that include the modified portion are marked
45 "invalid," signifying that the mapped tiles stored in the disk file no longer correspond to what is in memory. When it is desired to save the

changes, the changed tile blocks are decomposed as described above with reference to Fig. 2.

In a specific embodiment, the newly calculated tile blocks are stored at the end of the file, and the offsets and lengths for these tile blocks, as stored in the index, are updated to reflect the fact that these
5 tile blocks are now in new locations in the file and are of different lengths. The old versions of these tile blocks remain in the file, but the index contains no pointers to them. This allows the data to be saved without having to rebuild the file on every save operation.

This has the advantage of speed, and further does not
10 overwrite data during the save operation. Thus, it is potentially safer in the event of a computer malfunction. On the other hand, this embodiment wastes disk space. This can be addressed by rebuilding the file in response to a command. Alternatively, the problem of wasted disk
15 space can be alleviated by keeping track in memory of these empty spaces, and writing a tile block to the end of the file only if there are no empty spaces large enough to accommodate the tile block to be written. For safety, the tile block's old location, even if large enough, could be reserved until after the tile block has been safely written at a new
20 location.

Memory is typically a scarce resource when dealing with large images. For performance reasons, when retrieving tile blocks from the file it is preferred to maintain the compressed data in memory where possible. Additionally, the index in memory may advantageously store
25 additional information about the location of the tile blocks, such as whether they are in memory.

Specific Image Storage Format Overview

The section following this section contains a listing of a
30 source file Aff.h, written in the C computer language, showing a representative active file format according to the invention. While the source code is commented and is relatively self-explanatory, a few comments are in order. The so-called Magic number allows operating system utilities to quickly identify file contents, but has no additional
35 significance to the file format. The version field allows future expanded definition for the active file format.

The directory is stored at an offset rather than directly
after the file header so that a new directory can be written at the end of file and be made active by only overwriting the directory offset in the
40 header, rather than having to overwrite the old directory. As currently defined, fColorSpace in AFF.h is only exemplary, and illustrates only a few of the myriad of possible color spaces.

As can be seen, the actual filter coefficients are stored,
providing extra flexibility. The fReserved fields are for future
45 expansion as well as to make sure that the structs are powers of 2 in size so that index computations are efficient.

The basic iterative process for reconstructing an image to a desired resolution was described with reference to Fig. 4. The following discussion describes the method with specific reference to the source code. The steps outlined below are typical for reconstructing a given

- 5 desired layer (resolution level of the image):
 - 1) Read header and check if correct.
 - 2) Extract directory offset from header and seek to that position.
 - 10 3) Read fixed size AFFFFileDirectory structure.
 - 4) Set i = 0.
 - 15 5) Find data for the thumbnail by accessing its LayerImage description found at offset AFFFFileDirectory.fLayers[0].
 - 6) Load an operand buffer with the thumbnail image from its tiles found through information in the current LayerImage.fIndexOffset[kLLIndex].
 - 20 7) Use the information in the current LayerImage to locate the LH, HL and HH images. Combine this data with the image in the operand buffer and the given filter coefficients (in the current LayerImage) to reconstruct in an output buffer the image data for the next (i + 1) layer.
 - 25 8) Increment i by 1.
 - 30 9) If i = desired layer go to step 11.
 - 10) Find data for the next LayerImage at offset AFFFFileDirectory.fLayers[i] and read this data to make it the current LayerImage. Make the output buffer the operand buffer and go to step 7.
 - 35 11) Done.

Source Code for Specific Image Storage Format

```
#ifndef _AFF_H
#define _AFF_H
/*****
 *   File:                AFF.h
 *
 *   Function:            Definitions for Active File Format.
 *
 *   Author(s):           Dolf Starreveld (ds)
 *
 *   Copyright:           Copyright (c) 1992-1995 Storm Software, Inc.
 *                       All Rights Reserved.
 *
 *   Source:              Original.
 *
 *   Notes:
 *
 *   Review History:
 *
 *   Change History:
 *   95_10_22_ds          Added this header.
 *
 *****/
/* Some supporting type definitions */
```



```

typedef signed    short SInt16;
typedef unsigned  short UInt16;
typedef signed    long  SInt32;
typedef unsigned  long  UInt32;

typedef UInt32    FileOffset; /* Offset into file */
typedef UInt32    DataLength; /* Length of data in file */
typedef SInt32     Fixed16;    /* Fixed point with 16 bits fraction */

#define AFF_MAGIC 0x41464630 /* 'AFF0' */
#define AFF_VERSION 0x0100 /* Version 1.0.0 */

/*
 * This header is always found at the start of each AFF file.
 */
typedef struct {
    UInt32    fMagic;          /* Magic number */
    UInt16    fVersion;        /* AFF version used for this file */
    UInt16    fFiller;         /* Unused */
    FileOffset fDirOffset;      /* Offset to current directory */
    UInt32    fReserved[1];    /* Reserved for future expansion */
} AFFHeader;

/* This AFF directory provides information about the image as well
 * as information needed to locate all sub-components needed to
 * reconstruct the image at a desired resolution.
 * In this struct, dimensions and resolution refer to the original
 * image (in other words the largest reconstructible image).
 * This struct is really of variable length because the fLayers array is
 * really of a size determined by fNumLayers.
 * fLLayer contains the offset to the LayerImage structure for the
 * lossless completion layer, if present. Otherwise it must be 0.
 */
typedef struct {
    UInt32    fWidth;          /* Width in pixels */
    UInt32    fHeight;         /* Height in pixels */
    Fixed16    fHRes;          /* Horizontal resolution in ppi */
    Fixed16    fVRes;          /* Vertical resolution in ppi */
    UInt8     fColorSpace;     /* Image's color space */
    UInt16     fNumLayers;      /* Number of reconstruction layers */
    FileOffset fLLayer;        /* Offset to lossless layer (LayerImage) */
    FileOffset fLayers[1];     /* Offset to layer(s) (LayerImage) */
} AFFFileDirectory;

/* The following color spaces are exemplary */
enum {
    kColorGray8 = 1,          /* 8-bit grayscale */
    kColorGray12,             /* 12-bit grayscale */
    kColorRGB8,               /* 8-bit RGB */
    kColorRGB12,              /* 12-bit RGB */
    kColorCMYK8,              /* 8-bit CMYK */
    /* etc. */
}

/* Every layer in the file format has four sub-images:
 * LL (low-low), LH (low-high), HL (high-low), and HH (high-high)
 * images. Typically, the LL image is not present, because it can be
 * reconstructed from the previous layer. In some cases, however, there
 * may be a benefit to having this image available directly and so it
 * might be present in the file. For the final stage thumbnail image
 * this sub-image is always present.
 * If the image is not present, its fIndexOffset field must be 0.
 */
enum {
    kLLIndex = 0,             /* Possible indices for fIndexOffset */

```

```

        kLHIndex,
        kHLIndex,
        kHHIndex
    };
    enum {
        /* Values for fCompression */
        kCompNone = 0,      /* No compression */
        kJPEG          /* JPEG compression */
        /* Up to 256 compression methods can be defined here */
    };

    typedef struct {
        UInt32    fWidth;      /* Width in pixels */
        UInt32    fHeight;     /* Height in pixels */
        UInt16    fTileSize;   /* Tile size in pixels */
        UInt16    fTilesPerRow; /* Number of tiles horizontally */
        UInt8     fCompression; /* Compression method used */
        UInt8     fFilterSize;  /* # of coefficients in filters */
        SInt16    fLowPass[8];  /* Coefficients for low-pass filter */
        SInt16    fHighPass[8]; /* Coefficients for high-pass filter */
        UInt8     fReserved[2]; /* Currently unused */
        FileOffset fIndexOffset[4]; /* Offset to tile index arrays */
    } LayerImage;

/* In the LayerImage struct, each fIndexOffset points to a location in the
 * file where the tile index is stored for an image. At that location there
 * will be an array of TileIndexEntry structs (see below). This array will
 * be a one dimensional array with a number of entries equal to:
 * fTilesPerRow * ((fHeight + fTileSize - 1) / fTileSize) * numComponents
 * numComponents is the number of image components as implied by the
 * fColorSpace of the image, e.g. for an RGB image it will be 3.
 *
 * The index is addressed by a one dimension index that is computed as a
 * function of:
 * - image component
 * - pixel coordinate (h & v)
 *
 * The index first stores all entries for the first row of tiles for the
 * first image component (i.e. image components are stored in a planar
 * model). Next come subsequent rows of tiles for that component.
 * Next comes the first row of tiles for the second color component etc.
 *
 * Consequently the baseIndex for a plane is computed:
 *     baseIndex[comp] = comp * (fTilesPerRow * tilesPerColumn)
 *
 * To find the index entry for a tile that contains a given point in the
 * image, use the Point2TileIndex macro below to compute the index within
 * the plane and add that to the baseIndex for the desired component.
 */

typedef struct {
    FileOffset fOffset;      /* Offset to tile data */
    DataLength fSize;        /* Size of data for this tile */
} TileIndexEntry;

#define Point2TileIndex(h,v,lip) (((h) / (lip)->fTileSize) + \
                                   ((v) / (lip)->fTileSize) * (lip)->fTilesPerRow)

#endif /* _AFF_H */

```

Conclusion

In summary, it can be seen that the present invention provides elegant and powerful techniques for managing bitmap images. The storage format of the present invention provides random-access reading and writing (at the tile level), and can supply pixels at any resolution or

color space. By allowing selective retrieval at the resolution needed for output (print or display), the invention consumes CPU and I/O resources that are proportional only to the data needed, not to the size of the original input image. Thus the user can trade speed for image quality at will, while retaining the ability to retrieve the image at maximum resolution. This is significant, since it makes it possible to download only the thumbnail and thus browse images over a modem.

While the above is a complete description of specific embodiments of the invention, various modifications, alternative constructions, and equivalents may be used.

One possible extension relates to storing different portions of the logical file at different locations such as in different physical files on the same storage device or on different storage devices. As described above, file storage system 25 in computer system 10 may include drives at various remote locations, and there is no fundamental need for all the tile blocks to be stored at a single location. For example, it may be desirable to store the smaller layers on a fast local disk drive for rapid access and the larger layers on a remote network drive. This would allow the user to reconstruct the image for display relatively rapidly, while the retrieval of the tile blocks for the higher-resolution layers, if necessary, could occur relatively transparently.

This modification would require only that the offsets in the index for locating the tile blocks would also include a path name. For the specific source code described above, the type `FileOffset` is defined as an unsigned 32-bit integer (`typedef UInt32`), and the offsets (`fDirOffset`, the `fLayers` array, and the `fIndexOffset` array) are defined as `UInt32`'s. To accommodate a pathname, the `FileOffset` type can be defined as a structure containing a character string for the path and a `UInt32` for the offset value. An example of C code for implementing this is as follows:

```
typedef struct {  
    char        fPath[512]; /* Path to root of file */  
    UInt32      fOffset;    /* Offset into file */  
} FileOffset;
```

The `UInt32` within this struct would then be used for computing the locations of tiles in the same manner as the offsets in the C code in the previous section.

Additionally, storing different layers as separate physical files facilitates a mechanism for controlling access to the higher-resolution layers. The higher-resolution layers could be encrypted or made subject to password control.

A further possible extension relates to whether the intermediate reduced images (LL images other than the thumbnail) are included as part of the image storage format. As discussed above, it is only necessary to store the thumbnail and the complementary images; images at the maximum and

intermediate resolutions can be reconstructed. The particular image storage format described above allows for any given layer to have the reduced image (LL image). One way to provide user flexibility is to allow the user to specify which, if any, intermediate reduced images should be stored so as to avoid having to reconstruct them from the thumbnail. It is also possible for the program that reconstructs the image (as shown in Figs. 4 and 5) to alert the user on the basis of the number of times that the image has been retrieved to a particular resolution.

Therefore, the above description should not be taken as limiting the scope of the invention as defined by the claims.

WHAT IS CLAIMED IS:

1 1. A method of converting a bitmap input image into an image
2 storage format, the method comprising the steps of:
3 (a) at each of a plurality of stages including a first stage
4 and a last stage, processing an operand image to produce a reduced image
5 and a set of additional pixel data;
6 (b) in the first stage, using the input image as the operand
7 image;
8 (c) in each stage but the first stage, using the reduced image
9 produced in the previous stage as the operand image, the reduced image
10 produced by the last stage being referred to as the final reduced image;
11 the processing at each stage including the steps of:
12 (i) subjecting the operand image to a reduction
13 operation to provide the reduced image, the reduced image including
14 low-frequency information from the operand image and having fewer
15 pixels than the operand image; and
16 (ii) subjecting the operand image to at least one
17 operation that is different from the reduction operation to generate
18 the set of additional pixel data so that the reduced image and the
19 set of additional pixel data are sufficient to reconstruct the
20 operand image to a desired degree of accuracy;
21 (c) subdividing at least some of the sets of additional pixel
22 data into tiles; and
23 (d) creating the image storage format, the image storage
24 format comprising a representation of the final reduced image, a
25 collection of representations of each of the tiles of each of the sets of
26 additional pixel data from the plurality of stages, and an index that
27 enables determination of the respective locations of the representation of
28 the final reduced image and the respective locations of the
29 representations in the collection.

1 2. The method of claim 1, and further comprising the steps
2 of:
3 subdividing a particular reduced image produced in at least
4 one stage other than the last stage into tiles; and
5 including representations of each of the tiles of the
6 particular reduced image in the image storage format.

1 3. The method of claim 1 wherein:
2 the reduced image at any given stage has half as many pixels
3 in each dimension as the operand image communicated to the given stage;
4 and
5 the set of additional pixel data at the given stage includes
6 three pixel arrays commensurate in size with the reduced image at the
7 given stage.

1 4. The method of claim 1 wherein the representation of each
2 of the tiles of the final reduced image is a compressed version of the
3 tile.

1 5. The method of claim 1 wherein the representation of each
2 of the tiles of each of the sets of additional pixel data is a compressed
3 version of the tile.

1 6. The method of claim 1, wherein the representation of at
2 least one tile is a version of the tile that has been compressed by a
3 lossy compression technique, and further comprising the steps of:

4 (g) reconstructing the input image using the final reduced
5 image and the sets of additional pixel data to provide a lossy
6 reconstructed input image;

7 (h) comparing the lossy reconstructed input image with the
8 input image;

9 (i) determining an incremental image that, when combined with
10 the lossy reconstructed input image, provides the input image
11 substantially without loss;

12 (j) compressing the incremental image; and

13 (k) including the compressed incremental image as part of the
14 image storage format.

1 7. The method of claim 1 wherein the tiles in the sets of
2 additional pixel data are all the same size.

1 8. The method of claim 1 wherein the tiles in at least one of
2 the sets of additional pixel data are a different size from the tiles in
3 another one of the sets of additional pixel data.

1 9. The method of claim 1 wherein:

2 the reduced image at any given stage has half as many pixels
3 in each dimension as the operand image communicated to the given stage;
4 and

5 the tiles for the set of additional pixel data in the given
6 stage have half as many pixels in each dimension as the tiles for the set
7 of additional pixel data associated with the operand image communicated to
8 the given stage.

1 10. The method of claim 1, wherein said step of subjecting
2 the operand image to a low-pass reduction operation comprises the steps
3 of:

4 subjecting the operand image to a low-pass filter in a first
5 direction followed by decimation to provide a first resultant image; and

6 subjecting the first resultant image to a second low-pass
7 filter in a second direction perpendicular to the first direction followed
8 by decimation to provide the reduced image.

1 11. The method of claim 1, wherein said step of subjecting
2 the operand image to at least one operation that is different from the
3 low-pass reduction operation comprises the steps of:

4 subjecting the operand image to a low-pass filter in a first
5 direction followed by decimation to provide a first resultant image; and

6 subjecting the first resultant image to a high-pass filter in
7 a second direction perpendicular to the first direction followed by
8 decimation to provide a portion of the set of additional pixel data.

1 12. A method of manipulating a bitmap image in a memory of a
2 computer, the image being divided into tiles, the image being available in
3 the form of an image storage format, said image storage format comprising
4 separable information for each tile and an index enabling determination of
5 the location of the separable information for each of the tiles in the
6 image, said separable information for each tile being sufficient to
7 reconstruct the pixel data of said each tile, the method comprising the
8 steps of:

9 (a) loading the index from the image storage format into the
10 memory to create an in-memory version of the index;

11 (b) creating a data structure in the memory, the data
12 structure having an entry for each tile in the image, the entry for each
13 tile including a valid bit, the valid bits of all tiles initially having a
14 first state;

15 (c) loading into the memory at least a portion of the image
16 storage format, said portion corresponding to a selected plurality of
17 tiles;

18 (d) in response to user input specifying the modification of
19 one or more particular tiles in the image, setting the valid bit of each
20 of said one or more particular tiles to a second state different from the
21 first state; and

22 (e) in response to signals specifying a save operation,
23 determine the data structure those tiles for which the valid bit is in the
24 second state, and for each such tile for which the valid bit is in the
25 second state,

26 (i) updating the image storage format to include the
27 separable information of said such tile,

28 (ii) updating the in-memory version of the index to
29 reflect the new location for the separable information of said such
30 tile, and

31 (iii) updating the index in the image storage format to
32 reflect the new location for the separable information of said such
33 tile.

1 13. The method of claim 12 wherein:
2 the separable information for a given tile includes sufficient
3 information to enable the given tile to be reconstructed at any resolution
4 of a plurality of resolutions.

1 14. The method of claim 13 wherein:
2 the separable file information for the given tile includes a
3 thumbnail version of the given tile and one or more high-pass filtered
4 complementary images of the given tile, said complementary images being at
5 different resolutions of said plurality of resolutions so that the given
6 tile can be reconstructed at said any resolution by combining the
7 thumbnail version and at least a selected one of the complementary images.

1 15. The method of claim 12 wherein the data structure is a
2 table.

1 16. The method of claim 12 wherein:
2 said image is stored in a file formatted in said image storage
3 format;
4 said step (a) retrieves the index from the file;
5 said step (c) retrieves said at least a portion of the image
6 storage format from the file;
7 said step (e)(i) stores the separable information to the file;
8 and said step (e)(iii) updates the index in the file.

1 17. The method of claim 12 wherein said step (f)(ii) stores
2 the separable information to the end of the file.

1 18. A method of retrieving, at a desired resolution, a
2 desired portion of a bitmap image stored in an image storage format, the
3 method comprising the steps of:
4 (a) accessing a particular portion of the image storage format
5 to determine locations of representations of tiles of the image at a first
6 resolution that is lower than the desired resolution, which tiles
7 correspond to the desired portion of the bitmap image;
8 (b) accessing the particular portion of the image storage
9 format to determine locations of additional pixel data at the first
10 resolution;
11 (c) retrieving the representations of tiles from locations
12 determined in said step (a);
13 (d) retrieving the additional pixel data from locations
14 determined in said step (b);
15 (e) combining the representations of tiles retrieved in said
16 step (c) with the additional pixel data retrieved in said step (d) to
17 generate the desired portion of the image at a second resolution that is
18 higher than the first resolution; and

19 (f) if the second resolution is lower than the desired
20 resolution, repeating said steps (b), (d), and (e) with the second
21 resolution being treated as the first resolution.

1 19. A method of retrieving, at a desired resolution, a desired
2 portion of a bitmap image stored in an image storage format, the method
3 comprising the steps of:

4 (a) determining whether the image storage format contains the
5 image at the desired resolution;

6 (b) if so,

7 (i) accessing a particular portion of the image storage
8 format to determine the locations of representations of tiles of the image
9 at the desired resolution, and

10 (ii) retrieving the representations of tiles of the
11 image at the desired resolution, which tiles correspond to the desired
12 portion of the bitmap image; and

13 (c) if not,

14 (i) accessing the particular portion of the image
15 storage format to determine locations of representations of tiles of
16 the image at a first resolution that is lower than the desired
17 resolution, which tiles correspond to the desired portion of the
18 bitmap image,

19 (ii) accessing the particular portion of the image
20 storage format to determine locations of additional pixel data at
21 the first resolution,

22 (iii) retrieving the representations of tiles from
23 locations determined in said step (c)(i),

24 (iv) retrieving the additional pixel data from locations
25 determined in said step (c)(ii),

26 (v) combining the representations of tiles retrieved in
27 said step (c)(iii) with the additional pixel data retrieved in said
28 step (c)(iv) to generate the desired portion of the image at a
29 second resolution that is higher than the first resolution, and

30 (vi) if the second resolution is lower than the desired
31 resolution, repeating said steps c(ii), c(iv), and c(v) with the
32 second resolution being treated as the first resolution.

1 20. The method of claim 21 wherein said determining step
2 includes accessing a designated portion of the image storage format that
3 indicates which resolutions of the image are stored.

1 21. The method of claim 21 wherein the additional pixel data
2 are representations of tiles of complementary images at the resolution
3 that is lower than the desired resolution.

1 22. A method of storing a bitmap input image comprising the
2 steps of:
3 at each of a plurality of stages including a first stage and a
4 last stage, processing an operand image to produce a reduced image and a
5 set of additional pixel data;
6 in the first stage, using the input image as the operand
7 image;
8 in each stage but the first stage, using the reduced image
9 produced in the previous stage as the operand image, the reduced image
10 produced by the last stage being referred to as the final reduced image;
11 the processing at each stage including the steps of:
12 subjecting the operand image to a low-pass reduction
13 operation to provide a resulting image that includes low-frequency
14 information from the operand image and has a reduced number of
15 pixels, the resulting image being the reduced image; and
16 subjecting the operand image to at least one different
17 operation to generate the additional pixel data so that the reduced
18 image and the additional pixel data allow a desired degree of
19 reconstruction of the operand image;
20 subdividing the final reduced image into tiles;
21 subdividing each of the sets of additional pixel data into
22 tiles; and
23 creating a file that comprises a representation of each of the
24 tiles of the final reduced image, a representation of each of the tiles of
25 each of the sets of additional pixel data from the plurality of stages,
26 and an index that specifies the respective locations of the
27 representations of the tiles of the final reduced image and the respective
28 locations of the representations of the tiles of the sets of additional
29 pixel data in the file.

1 23. The method of claim 22, and further comprising the steps
2 of:
3 subdividing a particular reduced image produced in at least
4 one stage other than the last stage into tiles; and
5 including representations of each of the tiles of the
6 particular reduced image in the file.

1 24. The method of claim 22 wherein:
2 the reduced image at any given stage has half as many pixels
3 in each dimension as the operand image communicated to the given stage;
4 and
5 the set of additional pixel data at the given stage includes
6 three pixel arrays commensurate in size with the reduced image at the
7 given stage.

1 25. The method of claim 22 wherein the representation of each
2 of the tiles of the final reduced image is a compressed version of the
3 tile.

1 26. The method of claim 25 wherein each of the tiles of the
2 final reduced image is compressed by JPEG compression.

1 27. The method of claim 22 wherein the representation of each
2 of the tiles of each of the sets of additional pixel data is a compressed
3 version of the tile.

1 28. The method of claim 27 wherein each of the tiles of each
2 of the sets of additional pixel data is compressed by JPEG compression.

1 29. The method of claim 22, wherein the representation of at
2 least one tile is a lossily compressed version of the tile, and further
3 comprising the steps of:

4 reconstructing the input image using the final reduced image
5 and the sets of additional pixel data to provide a lossily reconstructed
6 input image;

7 comparing the lossily reconstructed input image with the input
8 image;

9 determining an incremental image that, when combined with the
10 lossily reconstructed input image, provides the input image substantially
11 without loss;

12 compressing the incremental image; and

13 including the compressed incremental image as part of the
14 file.

1 30. A method of manipulating a bitmap image in a memory, the
2 image being divided into tiles, the image being retrieved from a file that
3 comprises separable file information for each tile and an index specifying
4 the location of the separable file information for each of the tiles, the
5 method comprising the steps of:

6 retrieving the index into the memory;

7 creating a data structure in the memory having an entry for
8 each tile in the image, the entry for a given tile including a valid bit,
9 the valid bits initially having a first state;

10 retrieving at least a portion of the file corresponding to a
11 selected plurality of tiles;

12 converting the retrieved portion of the file to pixel
13 representations of the selected plurality of tiles in the image;

14 in response to user input specifying the modification of a
15 particular tile in the image, setting the valid bit for the particular
16 tile to a second state different from the first state; and

17 in response to signals specifying a file save, determining
18 from the data structure those tiles for which the valid bit is in the
19 second state, and for each such tile,
20 converting the pixel representation for that tile to a
21 file representation for that tile,
22 storing the file representation for that tile at the end
23 of the file,
24 updating the index in memory to reflect the new location
25 for the stored file representation, and
26 storing the updated index in the file.

1 31. The method of claim 30 wherein:
2 the separable file information for a given tile includes
3 information allowing the given tile to be reconstructed at any of a
4 plurality of resolutions.

1 32. The method of claim 31 wherein:
2 the separable file information for a given tile includes a
3 thumbnail version of the given tile and high-pass filtered complementary
4 images at a plurality of resolutions so that the given tile can be
5 reconstructed by combining the thumbnail version and selected ones of the
6 complementary images.

1 33. The method of claim 30 wherein the data structure is a
2 table.

1/3




FIG. 1 (PRIOR ART)

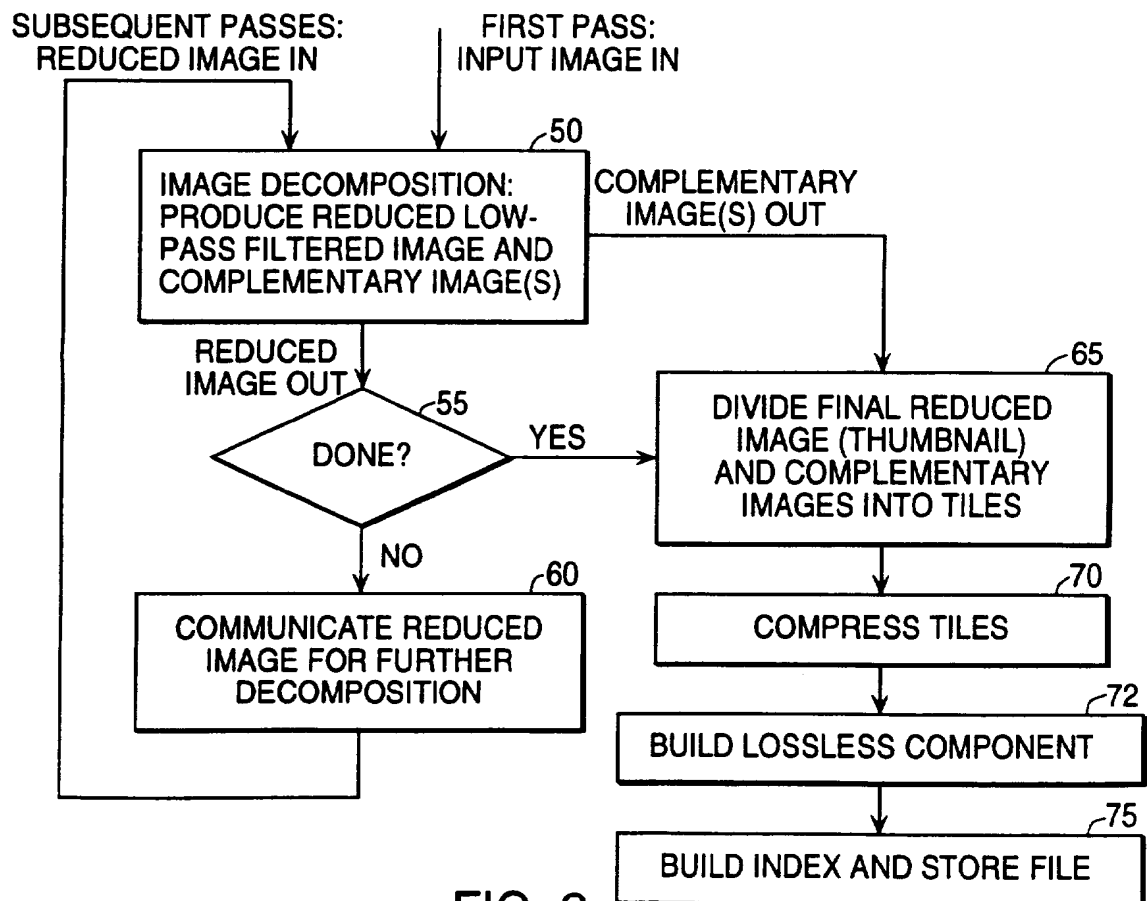
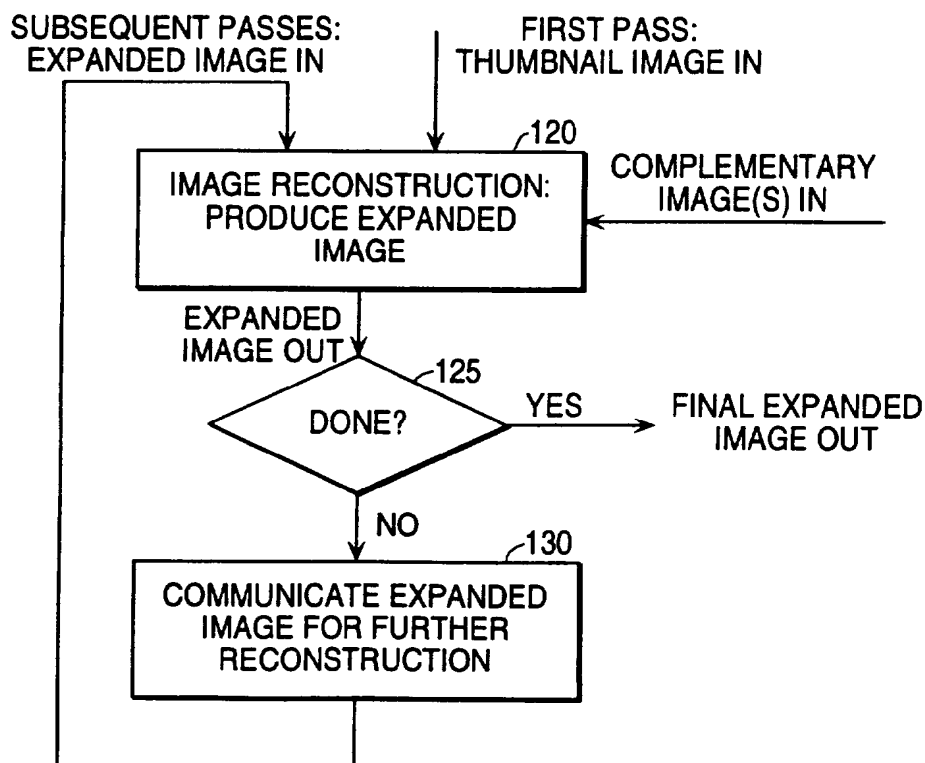
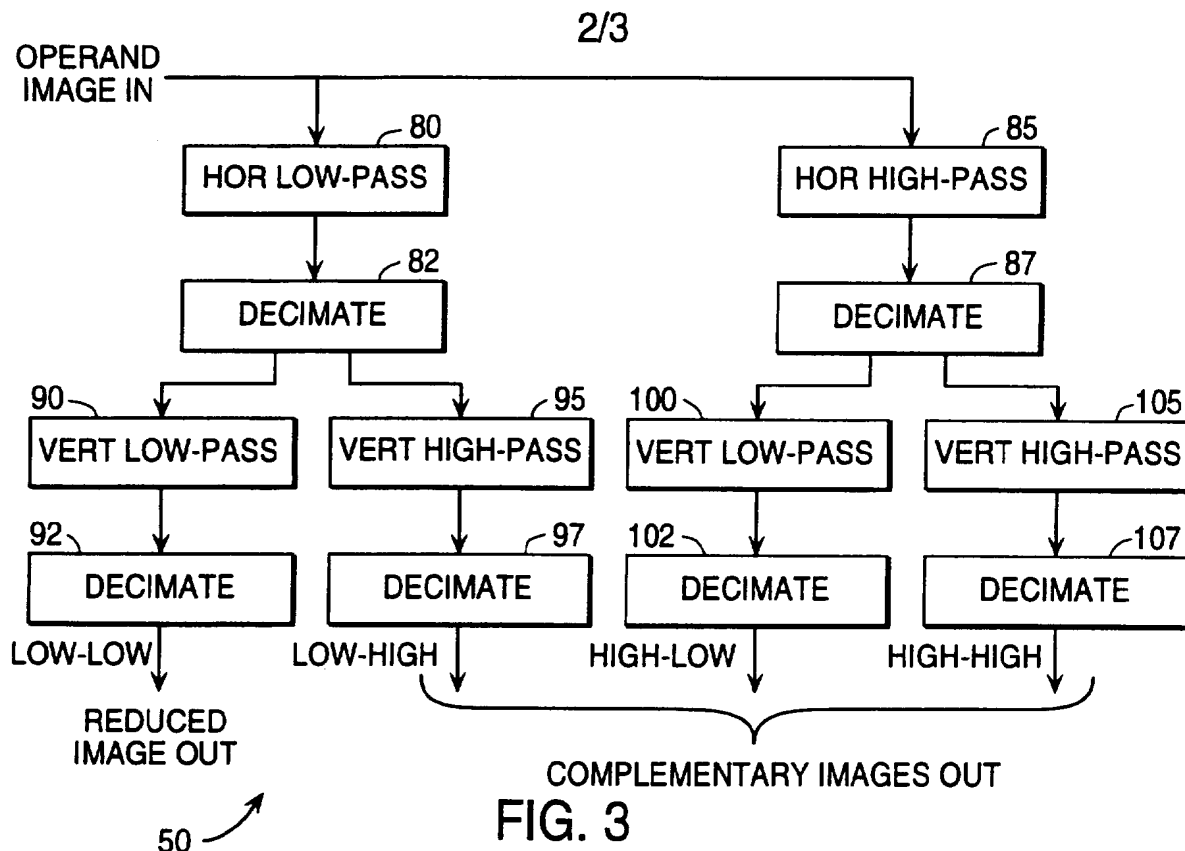


FIG. 2



3/3




FIG. 5

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US96/18017

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06K 9/36, 9/42, 9/54; G06F 15/00; H04N 7/12, 11/04

US CL : 382/232, 298, 302; 395/139; 348/384, 390

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 382/232, 298, 302; 395/139; 348/384, 390

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)
MAYA

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X -- Y	US, A, 5,048,111 (JONES ET AL) 10 September 1991, Figures 3, 5, 10, 13-14, 16, col. 7, lines 55-58.	1-9, 15-32 ----- 10-14, 33
X -- Y	US, A, 5,434,953 (BLOOMBERG) 18 July 1995, col. 10, lines 2-7.	1-9, 15-32 ----- 10-14, 33
Y	US, A, 5,333,212 (LIGTENBERG) 26 July 1994, abstract.	10-14, 33
A	US, A, 5,384,869 (WILKINSON ET AL) 24 January 1995, abstract.	1-33

☐ Further documents are listed in the continuation of Box C. ☐ See patent family annex.

* "A"	Special categories of cited documents: document defining the general state of the art which is not considered to be part of particular relevance	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"E"	earlier document published on or after the international filing date	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L"	document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O"	document referring to an oral disclosure, use, exhibition or other means	"Z"	document member of the same patent family
"P"	document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search
10 JANUARY 1997

Date of mailing of the international search report

12 FEB 1997

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Authorized officer

MONICA S. DAVIS

Facsimile No. (703) 305-3230

Telephone No. (703) 305-8576